



AFRL-AFOSR-UK-TR-2013-0046



ADAPTIVE PROBLEM SOLVING BY ANALOGY

Professor Georgi Petkov

**New Bulgarian University
Montevideo Str, 21
Sofia, 1618
Bulgaria**

EOARD Grant 10-3061

Report Date: July 2013

Final Report for 15 June 2010 to 14 June 2013

Distribution Statement A: Approved for public release distribution is unlimited.

**Air Force Research Laboratory
Air Force Office of Scientific Research
European Office of Aerospace Research and Development
Unit 4515 Box 14, APO AE 09421**

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<p>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 23 July 2013		2. REPORT TYPE Final Report		3. DATES COVERED (From – To) 15 June 2010 – 14 June 2013	
4. TITLE AND SUBTITLE ADAPTIVE PROBLEM SOLVING BY ANALOGY			5a. CONTRACT NUMBER FA8655-10-1-3061		
			5b. GRANT NUMBER Grant 10-3061		
			5c. PROGRAM ELEMENT NUMBER 61102F		
6. AUTHOR(S) Professor Georgi Petkov			5d. PROJECT NUMBER		
			5d. TASK NUMBER		
			5e. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) New Bulgarian University Montevideo Str, 21 Sofia, 1618 Bulgaria			8. PERFORMING ORGANIZATION REPORT NUMBER N/A		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) EOARD Unit 4515 APO AE 09421-4515			10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/AFOSR/IOE (EOARD)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-AFOSR-UK-TR-2013-0046		
12. DISTRIBUTION/AVAILABILITY STATEMENT Distribution A: Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The work on this project started on the basis of Associative Memory Based Reasoning (AMBR) model of analogical memory retrieval and reasoning. The goal was to explore the abilities of the model to do adaptive problem solving by making remote analogies, re-representing the target situation according to the current context, abstracting analogical cases and generalization of known solutions. In line with the project objectives, a number of new mechanisms were developed and their advantages were demonstrated in a series of simulations. In addition, some psychological experiments were conducted in order to verify the predictions of the model. Within the project, we also developed a set of technical tools for monitoring and controlling the operation of AMBR. The major results obtained were disseminated at conferences, workshops and other public events.					
15. SUBJECT TERMS EOARD, shock boundary layer interaction, Aerodynamics, Shock Waves					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 35	19a. NAME OF RESPONSIBLE PERSON Gregg Abate
a. REPORT UNCLAS	b. ABSTRACT UNCLAS	c. THIS PAGE UNCLAS			19b. TELEPHONE NUMBER (Include area code) +44 (0)1895 616021

Adaptive Problem Solving by Analogy

Effort sponsored by the Air Force Office of Scientific Research, Air Force Material Command, USAF

Grant number: FA8655-10-1-3061

Final Report

1. Overview

The work on the project started on the basis of Associative Memory Based Reasoning (AMBR) model of analogical memory retrieval and reasoning (Kokinov, 1994; Kokinov & Petrov, 2001). The goal was to explore the abilities of the model to do adaptive problem solving by making remote analogies, re-representing the target situation according to the current context, abstracting analogical cases and generalization of known solutions. In line with the project objectives, a number of new mechanisms were developed and their advantages were demonstrated in a series of simulations. In addition, some psychological experiments were conducted in order to verify the predictions of the model. Within the project, we also developed a set of technical tools for monitoring and controlling the operation of AMBR. The major results obtained were disseminated at conferences, workshops and other public events.

The next section of this document presents a brief description of the state of art of the AMBR of analogy-making. Section 3 described the new mechanisms which we introduced within the project: the anticipatory (transfer) mechanism, various ways for re-representation of the target, and the mechanism for knowledge abstraction. The anticipatory mechanism is the basis of the context-sensitive dynamic encoding of the novel situations, as well as of the analogical transfer of solutions from a retrieved base to the target problem. The re-representation of the target emerges from the interplay of the overlapping and dynamic mechanisms of AMBR. The initial sparse description of the target is dynamically enriched by relevant aspects of the context during the processes of retrieval, mapping, and transfer. The mechanism for abstraction lies at the core of processes of generalization both of the concepts and the situations.

Six series of simulations that highlight the most important accomplishments of the model are presented in Section 4. First, we simulate how remote analogies are made by means of a double analogy. The second simulation demonstrates the ability of the system to account for the processes of induction, deduction, and analogy by using the same mechanisms. The virtues of combining different reasoning modes are demonstrated by simulating how the two-sisters problem (two sisters quarrel for an orange and can't recognize that they want different parts of it and can divide it) can provide a solution to the problem of the Israel-Egypt conflict. The third simulation seeks to provide explanation of the well-known confirmation bias phenomenon. The fourth simulation demonstrates how the generalization of knowledge and the emergence of novel categories can result as a natural consequence of the process of analogy-making. Unlike other models, learning is not assumed as to be a separate mechanism but as part of the reasoning process. The fifth simulation deals with the context-sensitive and dynamic formation of the target representation. It demonstrates how the mechanism for abstraction may become essential

part of the representation of the current situation. The results of this simulation are complementary to the results from the first simulation. Both highlight dynamics and context-sensitivity of the target representation. Finally, the sixth simulation proposes a solution for the known binding problem within the AMBR modeling architecture. It accounts to how the mind binds together features, objects or other structures at varying levels of granularity in order to form meaningful wholes. AMBR's top-down mechanisms for transfer and anticipation-formation may substitute the complex calculation-expensive mechanisms that most of the models propose for solving this problem.

In Section 5, we describe some of the technical developments of AMBR: harmony tracking, visualization of the working memory, and a script language for setting-up and controlling simulations.

Summary and dissemination are presented in the last two sections.

2. Analogy-making and the AMBR modeling architecture

AMBR (Associative Memory Based Reasoning) is a model of analogy-making which is developed on top of the DUAL cognitive architecture (Kokinov, 1994). The building blocks of AMBR are hybrid nodes (micro agents) which exhibit both symbolic and connectionist properties. Each node has its localist meaning (an object, relation, scene, etc.), but at the same time it may be a part of the distributed representation of other nodes.

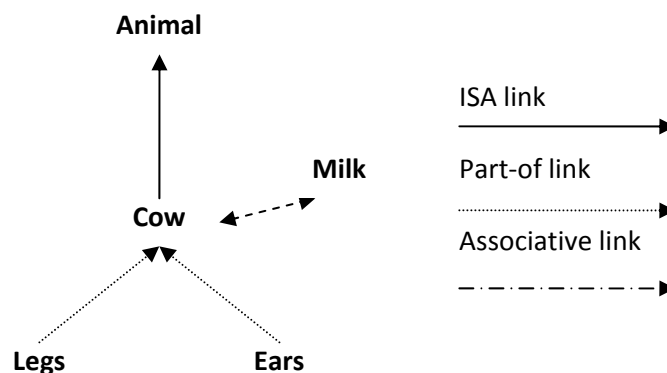


Figure 2.1. Knowledge representation in AMBR. There are three types of links between nodes.

The nodes are connected to each other by three types of connections - 'is-a', 'part-of' and 'associative' (Figure 2.1). 'Is-a' connections are used to represent conceptual hierarchy relationships, such as 'type-token' or 'class-subclass'. The role of 'part-of' links is to bring together elements which constitute a single entity, for example, all parts and properties of an object, or all elements of an event and the relationships between them (Figure 2.2). Associative links are used only for spreading activation, though other type of links can also spread activation.

There are two types of special nodes in AMBR – hypotheses and binding nodes. Hypothesis nodes represent mappings (analogical connections) between other nodes. Binding nodes are used to organize nodes into coalitions which collectively represent entities such as relations, events, concepts, episodes. The same binding nodes can participate in other coalitions and serve as the building block of distributed representation at a higher level.

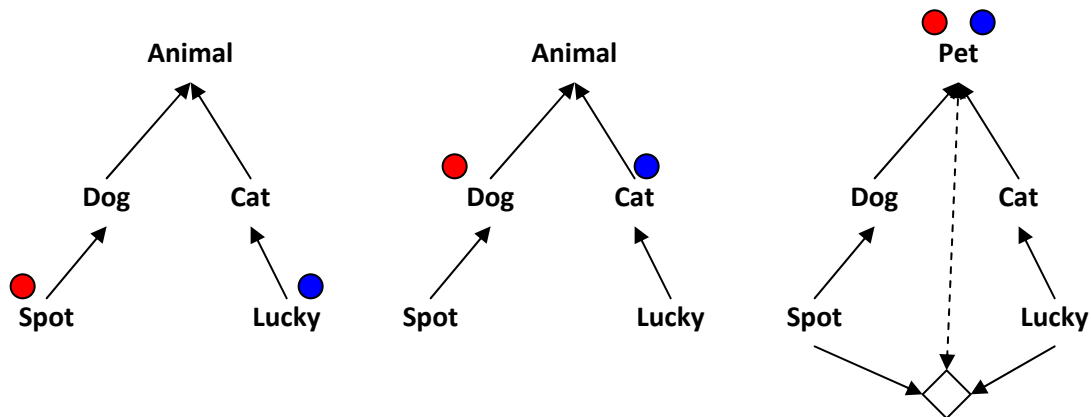


Figure 2.2. Examples of bindings nodes: a concept (to the left), an instance of a relation (middle) and script of an event (right).

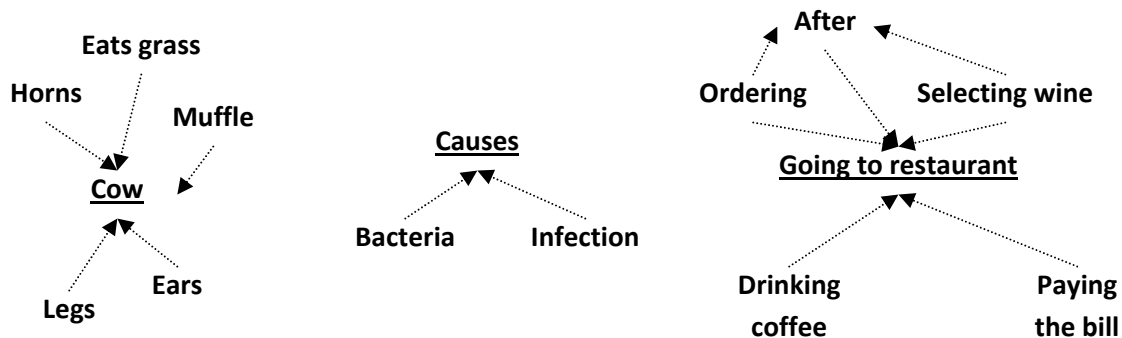


Figure 2.3. Establishing semantic similarity by marker passing. Instances of concepts send markers among the 'is-a' links. When two markers intersect at a node, a hypothesis node is created. It represents the mapping of nodes. The activation of the hypothesis node is a function of the activation of the nodes that were mapped, as well as the node at which their markers have intersected.

A hypothesis node is created when a semantic similarity is established. Semantic similarity is dynamically computed by a process of marker-passing (Figure 2.3). Hypothesis nodes can also be created due to top-down or bottom-up structural constraints (Holyoak & Thagard, 1989). The hypothesis nodes are created by using local information only – there is no central mechanism which monitors the whole network of nodes in AMBR. Hypothesis nodes which represent consistent mappings support each other by positively weighted associative links and inconsistent hypothesis inhibit each other (Figure 2.4). Thus the

hypothesis nodes create a constraint satisfaction network. The outcome of cognitive processes such as memory retrieval, recognition, categorization and analogy-making is determined by the resolution of the constraint satisfaction network.




Figure 2.4.Constraint satisfaction network - inhibition of rivaling hypothesis nodes.




Figure 2.5. Analogical transfer of binding nodes (anticipatory mechanism). The transferred anticipation-elements are drawn in red. The thick dotted line depicts the transfer process and is not part of the representation. Elements are transferred only to the target situation.

The role of binding nodes is to bind together elements of a distributed representation (Figure 2.2). Binding nodes are created in two ways – bottom-up and top-down. Bottom-up creation of binding nodes assumes that the perceptual input is already organized by automatic processes which lie outside of the scope of the AMBR model. Top-down creation of binding nodes is implemented by a process of

anticipation formation, called also transfer mechanism (Figure 2.5). Competing sources of analogical transfer can create binding nodes which are inconsistent. Inconsistent binding nodes start to inhibit each other and also form constraint satisfaction networks. In other words, there could be alternative interpretations of the perceptual input which compete with each other. We have used these mechanisms to model the recognition of ambiguous pictures and generate a non-trivial prediction, for which empirical support was found (Kokinov, Vankov, & Bliznashki, 2009).

The input to an AMBR simulation consists of a number of nodes which are connected to a source of permanent activation. The input nodes are also marked as targets so that the system tries to map them to other nodes retrieved from memory. The set of input and target nodes can be changed in the course of the simulation.

3. Extending AMBR - new reasoning and knowledge representation mechanisms

Analogical transfer

The transfer in AMBR emerges from the work of the anticipatory mechanism. If it happens all arguments of a certain relation from a base episode to be mapped to elements from the target, then the respective relation is transferred from the base to the target. However, the new relation is considered *as anticipation*. This is the basic idea of the anticipatory mechanism. It works according to the following algorithm: Each relation-agent that is from a retrieved base, not from the target, sends a special symbolic message to all its arguments after entering in the WM and after sending its usual marker. The respective arguments keep these messages in their symbolic buffers and if it happens for them to be involved in a hypothesis, they send back to the relation a symbolic answer with their respective correspondence. When all arguments of a certain relation answer with such messages, a copy of the relation is created. However, the arguments of the new relation are the respective correspondent elements to the arguments of the original relation. The newborn agent is a temporary-agent from a new type - *:anticipation-agent*.

Of course, it may happen one and the same argument to participate in several hypotheses simultaneously and thus, several different anticipation-agents may be formed. Actually, during a typical run of a simulation with the anticipatory-mechanism, a lot of anticipation-agents emerge and influence the further behavior of the system.

After its creation, the anticipation-agents behave just as instance-agents. They emit markers. They may be involved in hypotheses. They participate in the constraint-satisfaction network. However, they may die if their activation drops below the threshold. In addition, they may die if a special mechanism for their fizzling (see below) is triggered. The role of the anticipation-agents is to represent knowledge which is anticipated by analogy to be presented in the environment but the system is always ready to reject it.

Thus, the transfer of the solution during analogical reasoning emerges after creation of huge number of anticipation-agents; their interconnection in the constraint satisfaction network; promotion of some of them and fizzling of other by the attentual mechanism; and propagation of the winners. The set of winning anticipation-agents may be viewed as a transferred knowledge after analogy-making task.

Dynamic re-representation of the target

The original version of AMBR (Kokinov & Petrov, 2001) assumes that the input to the system is fully structured - all the possible relations between the input elements are present. This assumption suggests that the existing memories can't affect the representation of the target situation. In fact, this is how all the leading models of analogy-making work (i.e. Falkenhainer et al, 1989; Holyoak, & Thagard, 1989; Hummel, & Holyoak, 1997). One of goal of the current project was to show that analogy-driven problem solving could benefit from starting with a sparser and more surface description of the target problem, which is gradually enriched and deepened as more and more analogical cases are retrieved from the long term memory.

Following the main principles of the model, all mechanisms run in parallel and overlap each other. Thus, mapping process influences perception; transfer mechanism influences mapping.

Re-representation of the target may emerge from highlighting novel relations of the target during the transfer mechanism: After the initial very sparse description of the target, some initial bases are retrieved from the memory and some initial mappings emerge. Then, the transfer mechanism plays its role. Many anticipations that are important core elements of the retrieved bases emerge in the target description. The system actively searches for their presence in the environment and thus, gradually, the representation of the scene is enriched.

Additional way for re-representation of the target was demonstrated by AMBR - via double-analogy (see simulation_1 below). Once the target is presented, initially superficial analogous bases are retrieved from memory. However, if the earlier superficial analogy does not fit well, it may serve for additional search for both another bases for analogy and new aspects of the target situation.

Abstraction, concept acquisition, generalization of problem solutions

One of the weak aspects of the existing models of analogy-making, including AMBR, is that their operation heavily relies on the existence of pre-determined knowledge bases. There are two kinds of knowledge in AMBR – episodic and semantic memories. The difference between the two is blurred, because they are encoded using the same representational primitives and the system cannot differentiate between them by itself. However, it is supposed that the kinds of knowledge have different origins. Episodic memories stand for particular experiences, while semantic memory is thought to represent conceptualization of the external world. One of the goals of this project was to demonstrate how an analogy-making system can evolve its conceptual network by itself.

There are already at least two attempts to model the acquisition of concepts by analogy-making – SEQL (Kuehne, Forbus, Gentner, & Quinn, 2000) and DORA (Doumas, Hummel, & Sandhofer, 2008) modes. Both models follow the same approach to building new concepts. It is assumed that in order to create a new concept one has to compare two situations in which it is inherently present. In other words, the system tries to find the analogous counterparts in the two situations. Then, an abstraction is made, which consists only of those elements which were present in both situations and which were found to be

analogous. The abstraction is later refined by removing those elements which are rarely mapped in new comparisons.

We followed a similar approach in developing the mechanisms which allow AMBR to extend its semantic knowledge base. The basic idea behind the abstraction mechanism is that hypotheses about correspondences are turned into permanent nodes which reside in the long-term memory (Figure 3.1). There are a number of advantages of this approach.




Figure 3.1. New abstractions are created by hypotheses nodes representing the mappings between elements of analogous situations. The hypotheses nodes represented by diamonds. The newly abstracted structure consists of the very same nodes which are created during the mapping process.

Other models of abstraction by analogy, such as SEQL and DORA, assume that the end result of the abstraction process is a new concept, such as 'a dog', 'an ambulance', 'a red square'. To this end, abstraction is only used when comparing particular objects. The abstraction mechanism in AMBR is more general. An abstraction is just a set of interrelated nodes. It does not necessary represent a well defined concept. The system can built abstractions of whole situation, consisting of many objects and relations. Our assumption is that particular concepts emerge as abstractions are gradually enriched and refined in the course of the operation of the system. In other words, in order to acquire the concept of a 'dog', one does not need to present a series of dogs to AMBR, force it compare them and abstract a new concept. Instead, the concept of 'dog' gradually emerges as the system experiences various situation involving dogs which it finds analogous.

The implementation of the new concept acquisition ability is seamlessly integrated into the existing AMBR mechanisms. We do not propose a new model for abstracting knowledge by analogy, but reuse the fundamental mechanics of AMBR (spreading of activation, hypotheses for correspondence, constraint satisfaction) for extending its functionality. Abstraction is just one more process which occurs in AMBR. It runs in parallel with other processes and is subject to dynamic context effects at any given moment of the model operation. This approach opens the door to modelling various interactions between perception, memory retrieval, structural alignment and abstraction.

4. Some Simulations, Highlighting the Main Achievements During the Project

Simulation 1. Making remote analogies and the role of double-analogy

All known models for analogy-making either do not account to retrieval process (i.e. demonstrate the mapping only), either rely on a retrieval based on superficial similarity. AMBR architecture is unique in its dynamicity. It does not separate the sub-processes of analogy-making but view them as overlapping. Thus, all sub-processes influence each other and as a consequence, the initial superficial mappings may trigger further retrievals. Thus, the relevant remote base may be activated step by step. This process has been modelled and has been demonstrated by the first simulation:




Figure 4.1.1. Initial mappings and the consequent retrieval of novel relevant aspects.

Let the goal of the system is to find an appropriate remote analogy to a suicidal terrorist act, made by a single terrorist; and if possible, to transfer additional knowledge or, even a proposal how to prevent further similar acts. One superficially similar base is the situation of a japans kamikaze during the war. The system easily activates this base and launches the analogy (Figure 4.1.1). However, this analogy is

not good and falls later on. The reason is that one vivid aspect of the kamikazes is their motivation (Figure 4.1.1). The death for the japans is not an end of everything but just another possible method for solving problems. The kamikaze makes his suicide for the prosperity of his country, emperor, and family.

Once activated, the motivational aspect of the kamikaze situations tries to map with its analog in the terrorist situation. Thus, the question about the deep psychological motivation of the terrorist's act "crosses the mind", i.e. the system activates it. However, the encoded knowledge about the terrorist's motivation is that he is a foreigner from several years; he has relative good educational and professional successes but he is not happy. He has never overcome the cultural differences; the guilty that he has left his country; and the nostalgia.

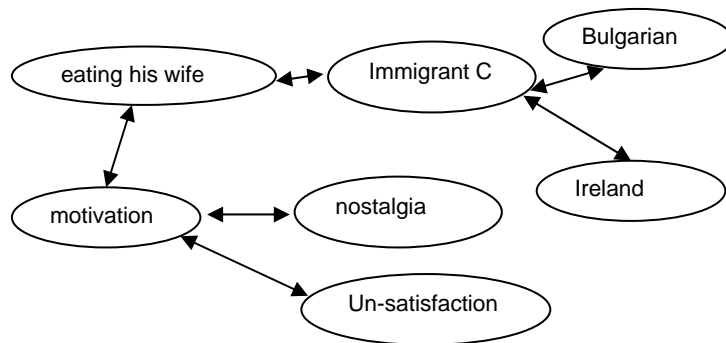


Figure 2.The representation of the 'Immigrant base'.

Once activated, this aspect of the target situation activates completely different base. Namely, the base of a Bulgarian emigrant in Ireland who has the same problems to adapt himself to a different culture and, as a consequence, he beats his wife (Figure 4.1.2). Nevertheless that this base seems quite different from the terrorist's one, it wins the analogy because of the deep structural analogy according to the motivation.

The last step for the system is to make a transfer. The story for the Bulgarian emigrant in Ireland has a happy continuation. This man has found a solution and has solved his problems. Actually, he has opened a Bulgarian restaurant and a small shop for traditional Bulgarian souvenirs. Thus, from one side, he has never uprooted fully from his country and, from other side, has deserved a respect from the Ireland people (Figure 4.1.3). According to the transfer mechanism in AMBR, however, the solution is not transferred directly but what is transferred is its super-class, according to the abstract hierarchy.

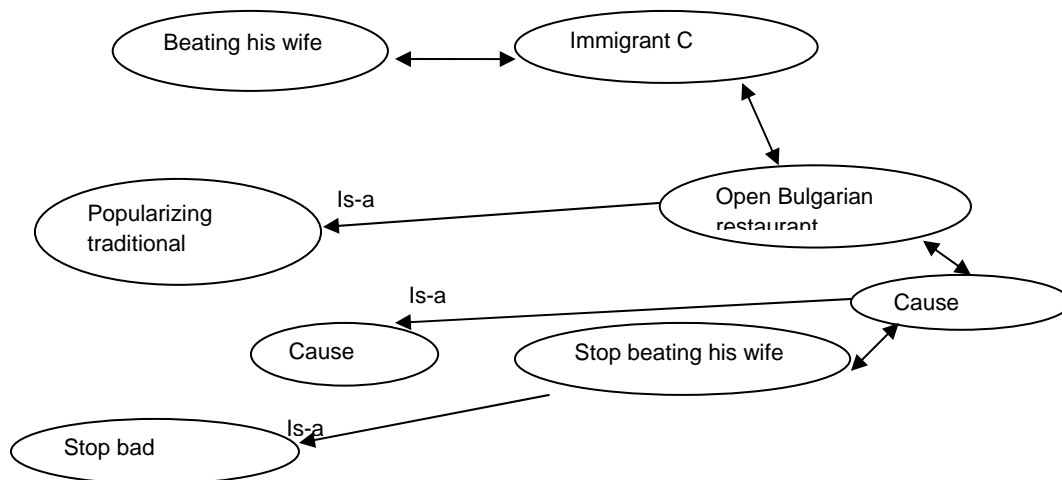


Figure 4.1.3. Representation of the ‘solution’ and its superficial level that is transferred.

Realizing this scenario, we demonstrated how the AMBR architecture can gradually re-interpret the target situation and how it can make remote analogies, activating the appropriate base step by step. The AMBR architecture is unique from this perspective. In addition, the ability of the architecture to make adaptive transfer has been demonstrated, as well as its goal dependent behavior, using both general knowledge and memorized situations for problem-solving.

Simulation 2. Unifying deductive, inductive and analogical reasoning

Traditionally, models of analogy-making have been isolated from and contrasted to models of deductive reasoning (Gentner, 1983, 1989; Holyoak & Thagard, 1989, Hummel & Holyoak, 1997). Somehow gradually analogy-making became a separate and important domain of study. However, may be deduction, induction, and analogy are not separate cognitive mechanisms, but rather slightly different manifestations of the same basic mechanisms (Kokinov, 1988) and the AMBR architecture can potentially account for this.

In order to demonstrate some of the AMBR’s important abilities in a coherent set of simulations, we decided to apply the model to a series of negotiation problems which require a trade-off solution (see Gentner, Loewenstein, Thompson, & Forbus, 2009).

An example of a trade-off problem is the classical story of the two sisters quarrelling over an orange (Figure 4.2.1) which is compared to the conflict between Egypt and Israel. Actually three simulations has been performed.




Figure 4.2.1. Representation of the two-sisters problem: two sisters quarrel for an orange and can't recognize that they want different parts of it and can divide it.

The first simulation models the mind of the first sister. Her goal is to make a shake, having an orange. Thus, there are two agents that receive initial activation: a representation of an orange is on the INPUT; a representation of a shake is on the GOAL. The activation spreads through the class hierarchy – to the concepts of ‘orange’ and ‘shake’; upward to the more abstract concepts; and then back to some of their instances. Relatively easily, the general knowledge of the recipe for shake is activated. There are other instances of orange and shake in the recipe and the mechanisms for marker-passing, the creation of hypotheses, and transfer do their job to produce the mapping between the given products and the recipe. Soon, the respective relations that are necessary for completion of the situation are transferred back from the recipe knowledge. Namely, the sister should take the orange; this implies that she can squeeze out the juice; this in turn is a necessary condition for making a shake. Thus, the system ‘solves’ the problem – the solution starts from taking the orange.




Figure 4.2.2. Part of the system’s representation of the situation for simulation 2 (only a part of the chain is shown).

The simulation successfully demonstrates the ability of the model to select from an un-separated general knowledge the relevant relations; to transfer them; and to combine them into a coherent solution. This differs from the traditional analogy-making tasks, in which the base situation is separated from the other knowledge.

The second simulation simulates the mind of a third person – a judge. There is again an orange on the INPUT, but an agent, which represents the relation that both sisters should be satisfied, is attached on the GOAL (Figure 4.2.2). The same long-term memory that has been used for the previous simulation is used.

The goal, however, is different and this changes dramatically the further representation of the situation by the system. It is easy for the model to activate the recipes for making shake and cake, and to transfer the respective relations. In other words, it can combine the two solutions from the previous simulations. However, this is not enough for achieving the goal. It is not possible for both sisters to take the orange.

Thus, no chain of relations to the goal is created and the activation continues to spread. Since both the juice and the peel are active, another piece of knowledge ‘springs up into the mind’ of the model. The juice, the peel, the seeds, etc. are all parts from an orange. Now knowledge of how to separate an orange into its parts becomes active. A different chain of transferred relations reaches the goal, wins the competition, and finally, a completely different representation of the situation has been made by the system (Figure 4.2.3).




Figure 4.2.3. The long-term memory has been enriched with a new base after the simulation 2.

The second simulation additionally demonstrated the ability of the system to use the basic AMBR mechanisms for solving problems that formally are not problems for analogy-making, but rather deductive tasks. Further, the simulation highlights the importance of the context-sensitivity of AMBR. Depending on the goal of the system, different relations may be transferred into the representation of the situation. The initial and the final representation of the situation may be viewed as two ends of a continuum of dynamic re-representations of the situation until the goal is reached.

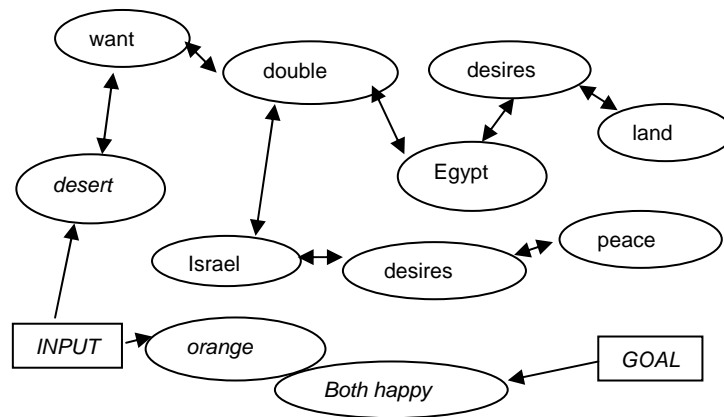


Figure 4.2.4. The initial state of the third simulation.

In the third simulation a representation of the classical Israel-Egypt problem is created and attached to the input of the system. The mind of a 'judge' is simulated. Thus, an instance of the relation 'both are satisfied' is attached to the GOAL. AMBR-agents for Israel and Egypt are attached to the INPUT (Figure 4.2.4). One instance of orange is also attached to the INPUT, simulating that the judge is by accident in front of a table with oranges on it. This is done to help the system retrieve the story about the two sisters. It is also hard for people to make such remote analogies (Gick & Holyoak, 1980). Maybe this is because it is difficult the respective remote bases to be activated. Maybe a certain non-trivial context is necessary in order the remote analogies to be initiated. The mechanisms of AMBR allow for this. Certain contexts may help it make remote analogies.

Egypt wants land and taking the dessert will satisfy it. Israel wants peace and taking the dessert will ensure it. All this knowledge is encoded in the long-term memory as general knowledge, analogically to the encoding of the sister's recipes. Simulating the point of view of Egypt (putting 'land' on the GOAL), the system would transfer the respective relations from the general knowledge and would conclude that it should take the desert. The same is for the Israeli point of view.

However, the desire that both of them be satisfied is on the GOAL list. The system cannot solve this problem by retrieving from general knowledge only. It cannot succeed in the same way as in the simulation 2, because there is no such general knowledge in LTM that land and peace are two separate properties of the dessert. Thus, the model makes an analogy between the target and the base already learned from the second simulation about how to divide the orange. Note that this analogy does not wait until the general knowledge is fully exhausted. Instead, everything runs in parallel. Of course, initially the contextual orange is mapped to the sister's orange, and the goal agent – 'both satisfied' to the base's goal. However, soon the pressure for consistency ensures the right mapping: Israel and Egypt to the sisters; and the dessert to the orange and so on.

The chain of relations to the goal is closed when the proposal to use separately the two properties of the dessert (it can be used to live on; it can be used as a buffer zone for ensuring peace, if it is demilitarized) are generated. We have not yet simulated how the transferred separation of the dessert properties may

be used for solving the task. i.e. how can the land be used for living by Egypt and at the same time be demilitarized for ensuring peace for Israel. This is part of our further work but it was already demonstrated that AMBR is able to combine relations from general knowledge in order to complete a representation of a certain situation.

Simulation 3. Modeling confirmation bias.

What happens if you make a wrong solution in a certain situation and later on face an analogical situation? Will the same (wrong) analogy will be used again, and if yes, will it affect your decisions and behavior? These questions are currently not addressed in the field of analogy-making. The vast majority of existing studies consider analogy-making only as a way of improving problem-solving and other cognitive skills, but not as a source of irrational behavior. One well studied phenomena of human's violations of the rational thinking is the so-called confirmation bias.

Confirmation bias is a robust behavioral effect in human thinking and decision making. The term confirmation bias combines various manifestations of the tendency people to search, retrieve, and interpret information in a way that confirms their current hypotheses. There are at least four different trends all called confirmation bias. First, people tend to retrieve from their memory more evidences than falsifications of their hypotheses (Cantor, 1979). Second, people tend to interpret the incoming information in a biased way that supports their current believes (Westen and co., 2006). Third, people tend to interpret memorized information in a biased way (Loftus, Palmer, 1974). Finally, people tend actively to search dominantly for information that confirms their current hypotheses (Wason, 1960). Thus, people dominantly tend to retrieve; interpret incoming; interpret retrieved; and actively search information that support their hypotheses. Although sometimes this is not the best strategy, the challenge in front of the cognitive models is to explain these psychological effects in terms of sub-processes that produce flexible and effective behavior in the most of the other cases.

The goal of the current simulation is to explain the confirmation bias in terms of analogy-making.




Figure 4.3.1. Full representation of the target situation. The model receives the three objects A, B, and C as initial input only. The relations should be anticipated and verified.

A single simple scene has been designed for recognition. It consists of three items (Figure 4.3.1) – three relations among them, and one relation from higher order (relation between relations). Only the three items A, B, and C are attached to the Input each run of the simulation and only they are a source of initial activation. There were not any goals presented to the system. The memory of the system consists of two encoded situations only (Figure 4.3.2). The first one is structurally more similar to the target, because the pressure from the higher-order relations usually is crucial. However, the second base also has its chances because share many elements with the target.




Figure 4.3.2. Representation of the two bases, encoded in the memory of the model

The initial activation spreads from A, B, and C to the respective concepts and then back to some memorized instances. The initial mappings between the target and base instances launch the transfer mechanism. If each one of two instances (e.g. A and B) finds a correspondence, then the active relations, involving A and B from the base should be transferred. For example, either R1 or R2, or both (if active) may be transferred as anticipations. Soon after the first anticipations emerge, a simulated attention system checks whether the most active anticipation is actually present in the scene. This has been modeled just by pre-defining in a list the presented relations. If the respective anticipation is not in the list, it is deleted. If it is, it becomes a regular part of the description of the scene.

In order to receive results that can be analyzed statistically, the simulation has been run 100 times on 100 different variants of the memory bases, varying randomly the strength of the links among the agents. The exact time when the anticipation agents for R1 and R2 are verified was encoded for each of the 100 runs. These anticipations are part of the target's description. After the verification, R1 will be rejected, R2 will remain. 53 times out of 100 the more interconnected anticipatory relation R2 has been verified (end recognized) firstly; 44 times R1 has been verified (and rejected) firstly, in 3 of the runs R1 has not been noticed at all. This result is a demonstration of the fact that the system retrieves in a biased way. It is important to note, however, that this priority was very tiny – 53 versus 44 times. Actually,

before ‘knowing’ what is presented in the scene, the system tries to verify many different relations. The only advantage for R2 comes from the fact that it has been involved in higher-order relation and as a consequence, it receives a little bit more activation.

At the end of the simulation, the activation levels of the two retrieved relations were recorded. These two agents are part of the system’s memory. The difference between the activation levels of the two retrieved relations was much more expressive: 90 times R2 was more active; R1 – only 7 times. This is an illustration of the biased interpretation of the retrieved information.

Finally, the activation levels of the two binding nodes (for Base 1 and Base 2) have been recorded at the end of each simulation. The activation level of the binding nodes indicate the degree of retrieval, hence the degree of usage, of the respective situation. 97 times Base 1 was the more active one, and only 3 times – Base 2. This result is in support of the ability of the model to interpret even the incoming information in a biased way.

Summarizing, the simulations models the human’s confirmation bias effects as a natural consequence of the AMBR’s basic mechanisms.

Simulation 4. Analogy-driven abstractions and the accumulation of knowledge

Much of the existing research on computational mechanisms of analogy-making has been concentrated on modeling single analogies, i.e. the mapping between a target situation and a single base which is retrieved from memory. However it is often the case that a given problem is analogous to several previous situations. The efficiency of a problem solving system will be greatly enhanced if it is able to benefit from accumulating more and more cases which describe analogous problems and their solutions. We have shown in the previous phase of the project how AMBR can benefit from making multiple analogies by using a superficially similar but structurally inconsistent base as a ‘bridge’ to finding a better solution to the target problem (Petkov, Vankov, & Kokinov, submitted). There is however an inherent problem with making multiple analogies, which can seriously affect the performance of analogy-making driven problem solver. Ideally, the system should be able to identify and solve a new problem better if it can retrieve more than one similar case. For example, if the task is to learn to recognize dogs, the system should perform better and better as it is given more examples of dogs. This is exactly what various neural network models of cognition achieve by adjusting their weights to correctly categorize each new example as it comes (not that they don’t have problems in doing this, such as catastrophic forgetting, French, 1999). However, merely supplying more examples of the same problem situation to analogical-making model of cognition, such as AMBR, can actually deteriorate its problem solving ability, rather than enhance it. Figure 4.4.1 illustrates the problem of having multiple bases which match the target situation.




Figure 4.4.1. The problem with multiple bases matching the target. The animal at the bottom is the target. The task is to categorize it. It matches the three bases. Three hypotheses for correspondence are created and they are competing with each other (the red lines stand for inhibitory connections). Whatever mapping finally wins, it will be evaluated as a relatively weak one because of the competition.

According to the definition of analogy-making (e.g. Kokinov & French, 2002), the target should be mapped to just one base. If more than one base is retrieved from memory, they will compete with each other, even though they might be analogous to each other. As a result, it will take AMBR more time to find the best mapping and the solution will be evaluated as worse compared to the case when fewer analogous cases are retrieved (Figure 4.4.2). Such a behavior is, of course, undesirable as it will not allow a problem solver to improve by building up knowledge. The problem can be solved by letting the system make generalizations of analogous cases. Each time a new case is given to the system, it searches for the best match already in memory and stores not only the new case, but also the generalization of the new case and the match. What is specific to our implementation of this idea within AMBR is that episodic and semantic knowledge (the cases themselves and the generalizations, respectively) is stored within the same media (AMBR long term memory) and used in the same way in future operations. In other words, a new target can be mapped either to a specific case which the system encountered previously or to a generalization between several such cases. Thus the system neither suffers from catastrophic forgetting (as PDP-like neural networks which store generalizations only), nor from undesired competition between multiple bases. Figure 4.4.3 shows how AMBR performance improves by making generalizations as it learns new examples of a certain situation.




Figure 4.4.2. The goodness of mapping as a function of the number of bases which match the target situation. The goodness of mapping is equal to the activation of the winning hypothesis for correspondence, the maximum value is 1. The goodness of mapping quickly decreases as more bases add to the long term memory, although they are all analogous to each other and to the target.




Figure 4.4.3. The goodness of mapping depending on whether generalizations are created. The goodness of mapping improves when adding more analogous bases, but only allowed to abstract structure from them.

Simulation 5. The role of generalizations in representing novel situations.

One of the main principles of the AMBR model of analogy-making is that all cognitive processes are continuous, run in parallel and interact with each other. The representation of the target problem is also a process which execution overlaps in time with other processes, such as memory retrieval and mapping. Therefore, the contents and structure of the existing memory traces may affect the way newly experienced situations are encoded. A number of simulations have been developed in order to reveal the dynamics of knowledge representation in AMBR (e.g. Kokinov, Vankov, & Bliznashki, 2009; Petkov, Vankov, & Kokinov, 2012). In order to further develop this line of research, we ran a simulation which aimed to show that the representation of a novel situation can be affected by the presence of abstract knowledge. The results of the simulation show that the presence of abstracted structures in long term memory changes the representation of novel situations by forcing the system to encode new entities as more similar to commonly occurring patterns. This lets the system benefit, in terms of both speed and accuracy, from having multiple analogous cases of the same problem. Moreover, the effect of abstract knowledge on building new representations prevents the model from finding 'superficial' analogies, which tend to have high surface similarity to the target problem, but share little or no common relational structure with it. It is important that the effect of abstraction on representing new target problems will be more substantial as the size of the long term memory grows. The following paragraphs briefly describe the simulations.

The knowledge base of the simulation consisted of 4 episodes (B1, B2, D, T). Episodes B1 and B2 were analogous, they both contained instances of the relations R1 and R2, and a single instance of the higher order relation R3. Episode D was used as a superficially similar distracter. Episode T was the target - we were interested how it would be represented. Figure 4.5.1 represents the structure of the four episodes. The episodes were attached to the input of the model in the following order: B1, B2, D, T. The key manipulation was whether the system was allowed to generate and keep the generalization of episodes B1 and B2. We measured performance by analyzing the structure of the input after T was presented at the input. In particular, we were interested whether the elements of T were predominantly mapped to the elements of D or to the elements of B1, B2 (or their abstraction). We also checked the activation levels of binding nodes transferred from the different bases. Table 4.5.1 presents the results. It is obvious that when the system was allowed to generate and keep the abstraction of B1 and B2, then the target T was encoded as more similar to B1 and B2. However, when the abstraction of B1 and B2 was not formed, then the target T was represented as more similar to D. In other words, when there are no abstracted structures, the system is unable to take into account the prevailing tendencies in its knowledge when representing the novel situations. The generation of abstract structures however lets the system behave in the desired way by new encoding new experiences in terms of what it already knows. Such behavior would allow the system to benefit from accumulating knowledge – the more analogous examples of a certain problem it is given, the easier it would be to discover the same problem pattern in novel situations.

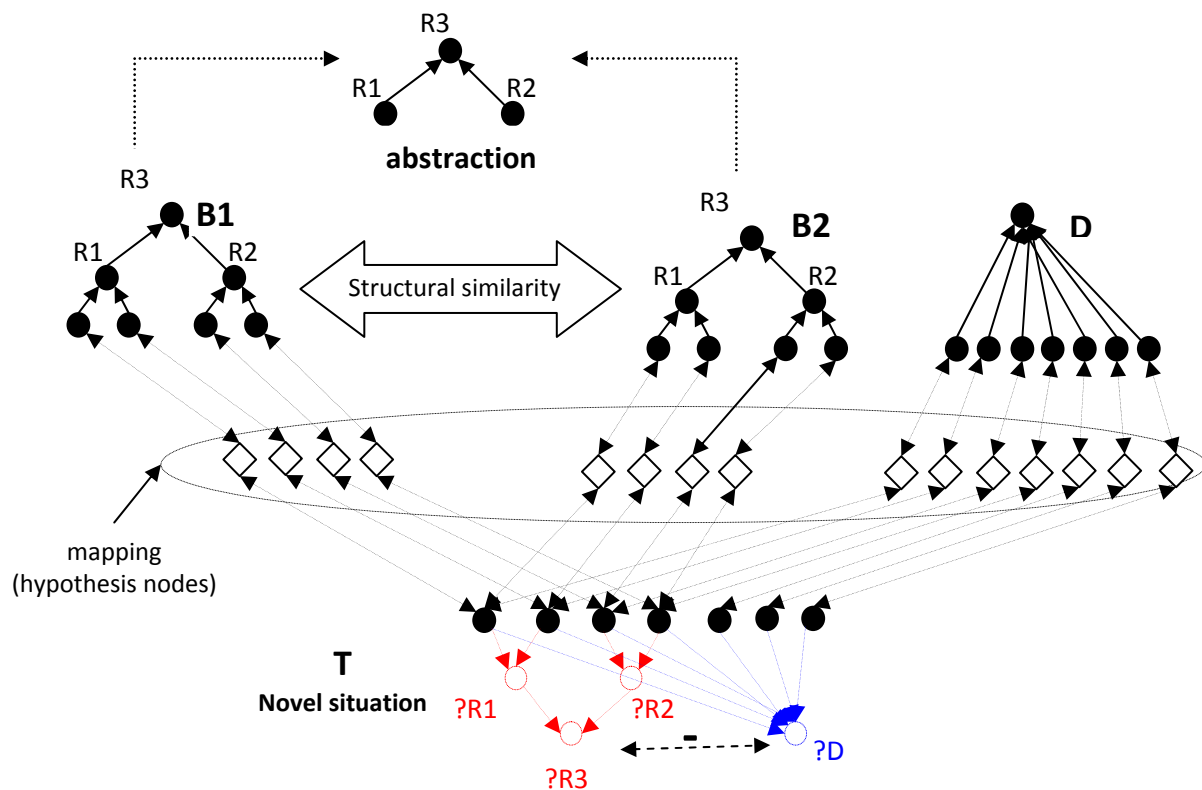


Figure 4.5.1. Overview of the memory structures used in the simulation. The target situation T was superficially similar to all the three bases B1, B and D. The binding nodes standing for the alternative representations of the target situation are given in red and blue. The activation levels of the competing binding nodes, which determine whether how T will be represented in long term memory, are given Table 4.5.1.

Binding node	Activation in abstraction mode	Activation when abstraction is allowed
?R1	0.853	0.176
?R2	0.853	0.175
?R3	0.990	0.698
?D	0.787	0.984

Table 4.5.1. Activation levels of the binding nodes representing the novel situation. When the system is allowed to abstract the common structure from analogous cases, then the abstracted structure is 'found' in the novel situation which is expressed by the higher activation levels of ?R1, ?R2 and ?R3 binding nodes. When no abstraction is allowed, then the high superficial similarity to the D episode overrides the structures transferred from the B1 and B2 episodes and results in a shallower representation of the target situation.

Simulation 6. The binding problem - how analogy could help disambiguate alternative representations of the target

The binding problem amounts to how the mind binds together features, objects or other structures at varying levels of granularity in order to form meaningful wholes (von der Malsburg, 1986). For example, if perceive a bit of fur, teeth and a wagging tail, how do we bind together these entities to form the percept of a dog. Or, how we bind and color and shape to perceive a blue square. The binding problem is one of the most pervasive ones in the cognitive sciences and has numerous instantiations, for example in language processing (how do we segment speech sounds into words) and vision (how do we determine whether a set of visual features belong to the same object). The binding problem is also relevant to problem solving as long as prior to finding a solution to a problem, one has to build its representation by finding the most appropriate organization of the target features. A particular example of a binding problem occurs when the relations between the objects in the target situation can be described in alternative ways. This is particularly important for models such as AMBR, which doesn't assume that all the relevant structures in the target situation are hard coded, but views representation building as intrinsic to the reasoning process.

Traditionally, solutions to the binding problem involve detecting regularities in the input, which can be used to infer which of its parts should be grouped together. For example, the gestalt psychologist developed a number of rules which prescribe when and how visual features are bound together. There is however a growing understanding that the solution of the binding problem should be sought in top-down processes (f.e. Lollo, 2012). The AMBR model provides a computational implementation of the idea and the grouping of elements in meaningful wholes is driven by background knowledge and top-down processes. In AMBR, there is a dedicated mechanism for binding things together – the transfer of binding nodes. For example, activating the features 'fur', 'teeth' and 'wagging tail' will activate the concept of a dog and the target features will map to the features participating in the concept representation. Then a binding node will be transferred from the dog concept to the target representation which will serve to bind together the features which belong together. It may happen that there are multiple binding nodes which are trying to organize the same target elements into different structures. If this happens, the ambiguity is resolved by a process of constraint satisfaction. The question which we tried to address during the last months is how effective this process is. To this end, we ran a series of simulations which involved building the representation of a target problem, which bears alternative interpretations. In one of these simulations, we checked whether AMBR tends to build structurally rich representations of the target problem. Figure 4.6.1 presents the elements of the long term memory, which was used in the simulation. The target situation consisted of four nodes, which could be grouped together in a number of alternative ways (Figure 4.6.2). There is particular way to bind together the objects which leads to building a structurally richer representation. Note that the initial activations of the nodes and the connection weights did not favor in any way the structural representation. The only reason to select it would be that it is structurally more systematic and thus provides more information about the target than the alternative representations. The simulation showed that indeed AMBR preferred the structurally richer representation of the target (Figure 4.6.3), even when one of the shallower solutions (Surface 2) was primed by making its relations more active.

Figure 4.6.3 also shows that the systems needs some time to build and select the structural representation of the target, which is consistent with psychological findings.

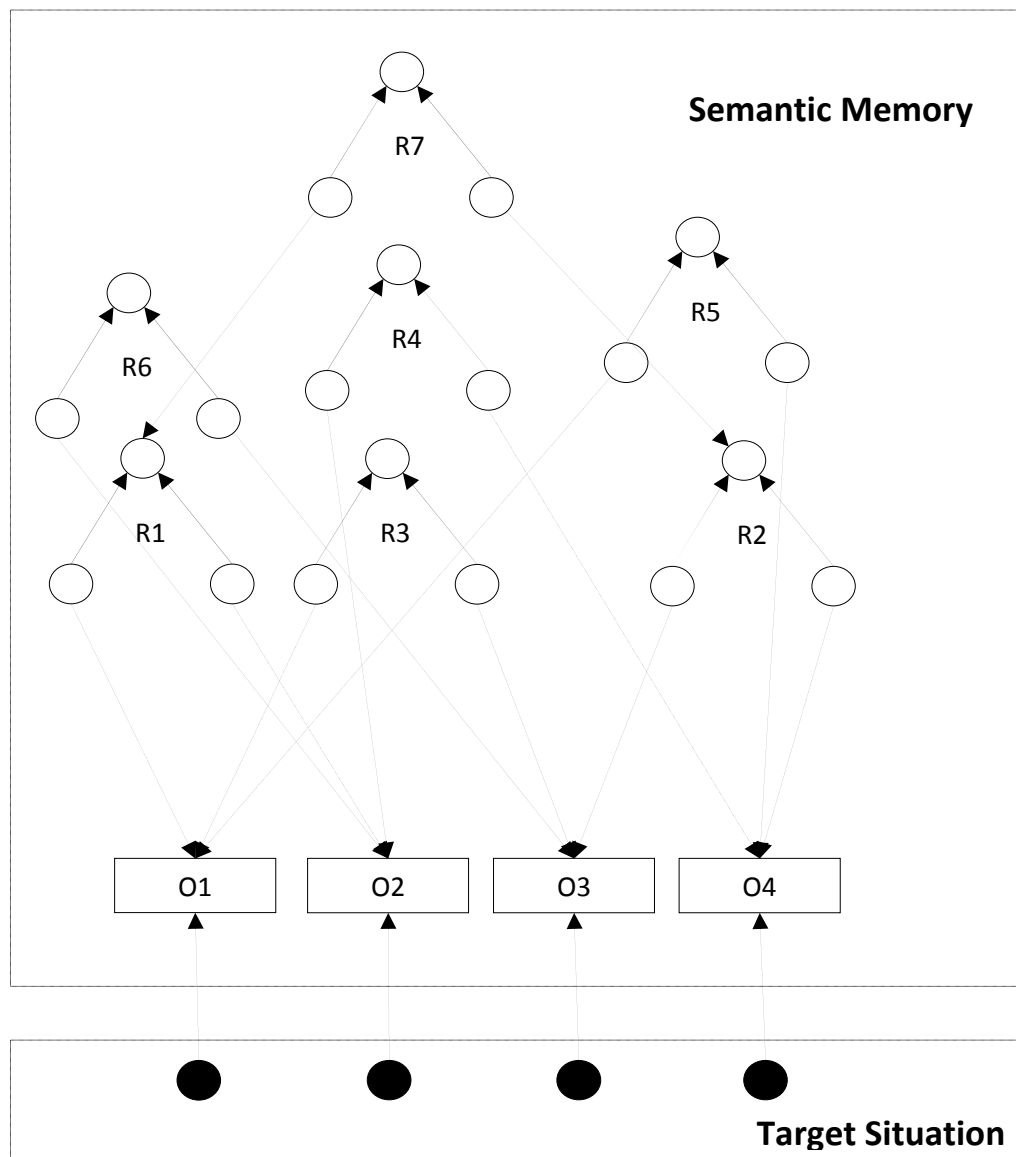


Figure 4.6.1. The memory structures taking part in the binding simulation. The dotted lines represent 'is-a' links and the solid lines stand for 'part-of' links. There were four objects in the target situation which were instances of four different types (O1-O4). There were also 6 different relations that were defined over the same set of objects (R1-R6) and a single higher level relation (R7) defined over R1 and R2. The binding of the target situation is ambiguous because all there are alternative ways to group the target objects by the lower level relations (see also Figure 4.6.2). Two of the lower-level relations (R5 and R6) were made a little bit more active than the rest.

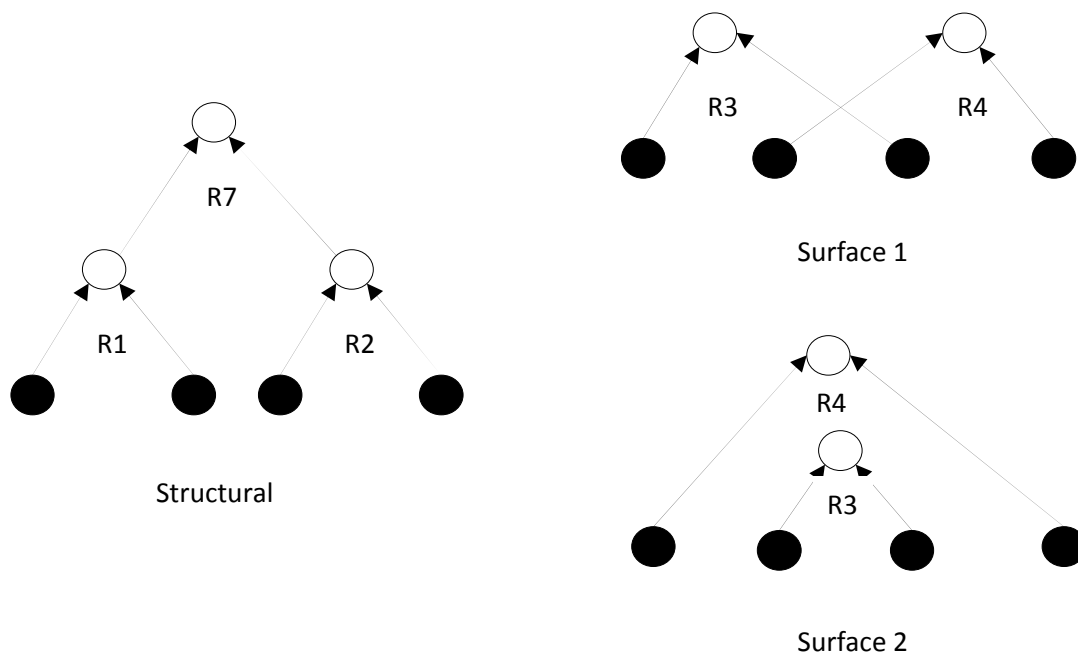


Figure 4.6.2. Three alternative bindings of the target situation. The target objects could be grouped in a number of alternative ways (not all possible bindings shown), depending on which of the lower-level relations are used. There is one particular binding which leads to discovering a higher level relation and thus results in building a 'deeper', structural representation.




Figure 4.6.3. The average activation of the alternative bindings of the objects (i.e. the average activation of the participating binding nodes) as a function of AMBR time. Initially, the structural representation had lower average activation, as the higher-level relation needed some time to be discovered. The difference between the two surface bindings was due to the slightly higher initial activation of the lower level relations taking part in the 'Surface 2' binding. In the course of time, the structural binding becomes the dominant one because of the AMBR

mechanisms for structural consistency. The simulation demonstrates AMBR's ability to represent the target situation in the most coherent, or structurally systematic, way.

The simulations of how AMBR resolves the binding problem have important implications for its problem solving performance. The core feature of our approach to problems is analogical reasoning. In order to make use of the power of analogical reasoning, one needs to provide with structurally rich representations. The simulation showed that our model is capable of building such representations even when more active structurally simpler alternative exist.

5. Technical Developments

Harmony tracking

One of the fundamental ideas underlying the AMBR model of analogy-making is that the process of mapping unfolds in time. This means that the mapping between a base and the target situation at a given moment may not be the best which the system is able to find if it is given more time. Generally, the system tends to evaluate superficially mappings higher at its early stages of operation and later on converge to 'deeper', more systematic structural solutions. It is therefore of vital importance to allow AMBR to run as long as it takes to find the best mapping. At the same time, a practical problem solver is expected to come up with a solution in a reasonable amount of time. In its original version, AMBR uses a locally defined rating mechanism as a criterion to stop the simulation and produce the final mapping. The idea of this mechanism is to keep track of which is the strongest mapping of target entity (in terms of its activation level) and terminate the simulation if it doesn't change for a certain number of time steps. The problem with this approach is that all mapping competitions must be resolved in order to finalize the simulation. If the constraint satisfaction network is big enough, it can take a lot of time to resolve all local competitions and the process is also computationally intensive. Moreover, sometimes it happens that the activations of two competing mappings oscillate periodically and the rating mechanism never comes to a solution. Another problem is that the rating mechanism has a parameter (the number of times required to find choose a 'winning' mapping) which is dependent on the network size - larger networks settled down slower. Therefore, a more general stopping criterion is needed. Holyoak and Thagard (1989) suggested that the status of a constraint satisfaction network at a given moment t can be described by calculating its harmony value $\sum w_{ij}a_i(t)a_j(t)$, where $a_i(t)$ is the activation of element i at time t and w_{ij} is the strength of the connection between elements i and j . We implemented the harmony function in the AMBR model and developed a new method for determining the moment in which the simulation can be stopped and the final solution produced. We also found out that we can use the dynamics of the harmony value to evaluate AMBR performance under different settings. For example, we found out that the harmony converges a bit faster but more reliably to an attractor when the system is allowed to make abstractions between analogous cases (Figure 5.1).




Figure 5.1. The dynamics of harmony depending on whether generalizations are created or not. The long term memory consisted of 6 analogous bases. When no generalizations were created the harmony function first converged to a local attractor and then found the better one. The presence of generalizations of analogous bases in the long term memory led to somewhat slower convergence, but there were no abrupt changes. The figure demonstrates how the harmony function can be used to keep track of the mapping processes and decide when to stop the simulation.

Visualization of AMBR working memory

One of the specifics of the AMBR model of analogy-making is that the knowledge is represented in a highly distributed manner. Even simple simulation may lead to the retrieval of tens of nodes from the long term memory and the formation of hundreds of temporary correspondence nodes. For example, Figure 1 contains about 50 nodes involved in the simulation and about the same number of other nodes is omitted for clarity. The abundance of various types of nodes and the connections sometimes makes it difficult to track the operation of the model. It is hardly possible to get a general impression of what is going on by looking at the list of active nodes (Figure 5.2). This is why we looked for a way to visualize the contents of the working memory of AMBR and track its operation. To this end, we decided to use the open source graph visualization toolkit Nodes3D (<http://brainmaps.org/index.php?p=desktop-apps-nodes3d>). Nodes3D has been designed for viewing complex neuroanatomical data, including large numbers of nodes with high connectivity. It provides a rich set of options for viewing graphs, including automatic scaling and arrangement, switching between 3D and 2D views, focusing on selected parts of sub graphs and others. In order to make use of the software, we developed a tool for encoding the

current contents of the AMBR working memory into a format which is readable by Nodes3D. Figure 5.3 shows the contents of the AMBR working memory looks like when viewed by Nodes 3D.

```

C:\WINDOWS\system32\cmd.exe
Y4<->R2-O2: 0.6923415797000000
Y4<->X4: 0.4504593582000000
?R1-2: 0.8530856945000000
?R2-2: 0.8530856945000000
Y5<->D5: 0.3522519518000000
Y6<->D6: 0.3522519518000000
Y7<->D7: 0.3522519518000000
?R1-2<->?R1-0: -0.7792918593000000
?R1-2<->?R1-1: 0.7001407879000000
?R2-2<->?R2-0: -0.7792918593000000
?R2-2<->?R2-1: 0.7001407879000000
?D-0: 0.7872482123000000
Y1<->A-X1_R1-O1: 0.7713248173000000
Y2<->A-X2_R1-O2: 0.7713248173000000
Y3<->A-X3_R2-O1: 0.7713248173000000
Y4<->A-X4_R2-O2: 0.7713248173000000
?R1-2<->A-?R1-1_R3-O1: 0.8864641610000000
?R2-2<->A-?R2-1_R3-O2: 0.8864641610000000
?R1-2<->R3-O1: 0.7733183384000000
?R2-2<->R3-O2: 0.7733183384000000
?R3-2: 0.9904748481000000
?R3-2<->?R3-0: 0.4119754336000000
?R3-2<->?R3-1: 0.6265917327000000
?R3-2<->A-?R3-1_?R3-0: 0.7339402970000000
Press any key to continue . . .

```

Figure 5.2. A snapshot of the output of an AMBR simulation. The list contains the contents of the working memory. There could be hundreds of active nodes and it is difficult to get a general impression of what is going on.

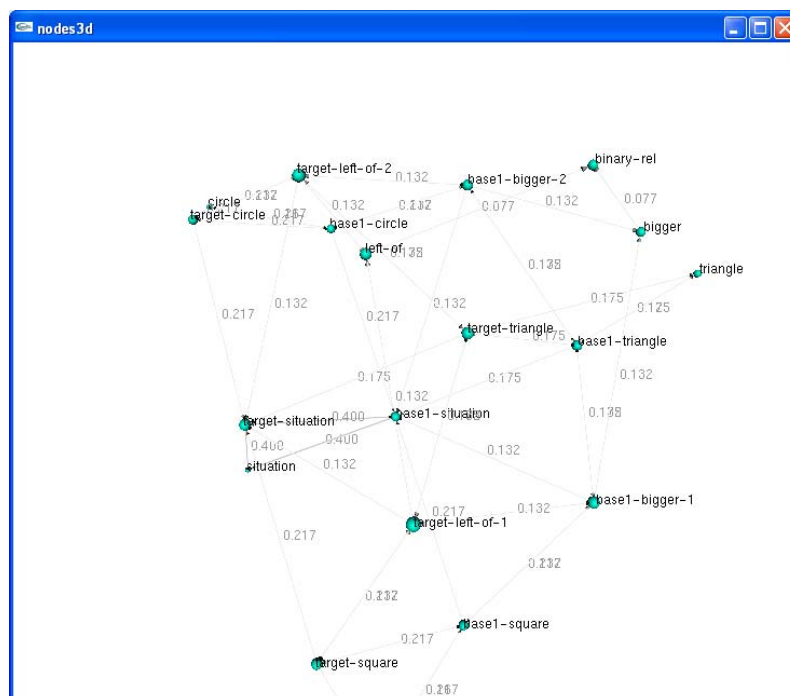


Figure 5.3. Visualization of the AMBR working memory. My using the mouse and the keyboard one can attend to specific parts of the network or zoom is out and remove unnecessary details in order to get a more general view.

Setting up AMBR simulations

The second main problem which we worked on during the first half of third year of the project concerned setting up AMBR simulations. We developed a kind of scripting language which allows us to easily run new simulations, track their dynamics and control them while they are executing. Figure 4 shows three fragments of the script, which we used to setup and run the simulation described above. In the first fragment, a new structure (a relation called R3) is added to the long term memory. The new structure consists of three nodes (R3, R3-O1, R3-O2). The arguments of the relation are connected to the relational binding node by 'part-of' links (denoted by '@') with the default weight of 1. The arguments of the relation are also connected to the O1 and O2 concepts by "is-a" links (">") with 0.5 weight. The new structure is tagged as a concept which allows it to be transferred to the target. The "#" sign is used to create associate links with a given weight. The second fragment shows how it is possible to track the execution of the simulation. The first line defines a moment in time (in AMBR time units) and the next two lines print the activation of two particular memory structures (anticipated binding nodes in the target situation). The third fragment shows how it is possible to print the contents of the whole memory (both LTM and STM), as well as run a command which controls the execution of the simulation (just terminates it in this particular example). There are a number of other commands and parameters which can be used to set up and dynamically control AMBR simulations. For example, it is possible to simulate interactive input to the system, by adding nodes to the target situation at certain moments in time. The new way of setting up AMBR simulations differs mainly in the ability to make change to the simulations while they run.

```
...

R3
R3-O1
R3-O2
R3-O1@R3
R3-O2@R3
R3-O1>O1
R3-O2>O3
O1#0.5#R3-O1
O3#0.5#R3-O2
/tag R3 CONCEPT

...

/time 10
/print act ?R1-0
/print act ?R3-0

...

/time 1000

/print memory

/terminate
```

Figure 5.4. A sample of a script defining and controlling an AMBR simulation.

Empirical validations of AMBR mechanism

The goals of this project did not suggest doing any empirical work, we ran a single experiment in order to test an intriguing prediction raised by Simulation 1. The goal of the experiment was to check whether a superficially similar base may serve to re-represent the target and serve as a 'bridge' towards finding a less accessible, however structurally consistent and overall more appropriate base. Participants were asked to read three stories and then to evaluate the similarity between the first and the last story. The manipulation was the contents of the second story. Half of the participants were given a second story which was superficially similar to the first one and could serve as a bridge analogy to the third one. The second story presented to the other half of the participants was also superficially similar to the first story but could not serve as a bridge analogy. The results indicated that those participants who were given a bridging analogy found the first and the third story more similar. More details about the experiment are given in the Appendix A.

6. Summary of the project achievements

The major achievements of the project are twofold. In the first place, we elaborated the existing reasoning and knowledge representation mechanisms of AMBR which lead to a better understanding of what computations underlie analogy-making. In particular, the building blocks of AMBR - spreading activation, marker passing, hypothesis formation, and constraint-satisfaction were amended with mechanism for generating anticipations. This set of basic mechanism was used to model the main sub-processes of analogy-making: representation of the target, retrieval from memory, mapping, transfer, and learning. The whole cycle from target representation through retrieval, mapping, transfer to learning was simulated for the first time during the APSA project. The principle that all mechanisms overlap and influence each other led to the development of novel views of how some of the sub-processes of analogy-making may be modeled: dynamic context-sensitive and goal-sensitive representation of the target; gradual mapping with superficial and structurally similar bases; emergent transfer of the solution on the basis of verification and rejection of goal-sensitive anticipations formation; learning as a natural effect of the process of reasoning and problem-solving. The integration of the old mechanisms of AMBR and the new ones leads to simulating processes of - context-sensitive encoding of new situations; adaptation of existing memory traces; generalization and abstract concept formation; generalization of known solutions.

In the second place, the project succeeded in showing that analogy-making can serve as the basis for adaptive problem-solving. In a series of simulations was demonstrated the ability of the system to make analogies between remote domains; to assimilate novel information in terms of what the system already 'knows'; to generalize cases and find abstract solutions which can be later applied to novel problems; to use abstract knowledge in order to enhance the retrieval of superficially dissimilar but relevant cases.

It is important to note that we conducted simulations at various levels of representation. Successful simulations were performed on the level of visual primitives and simple relations; on the level of mere

concrete and more abstract concepts; as well on the high level of analogical reasoning between remote domains.

Within the project, we also developed several tools which enable running and understanding complex simulations. The tool for harmony tracking may serve for critical analysis of the model, its evaluation and comparison with other models, including classical network mechanisms. The tools for visualizing the internal state of the working memory helps the modeler understand how the system solves the problem at hand.

Whereas the main objective of the project was achieved, new technical and scientific questions arose during the work on the project. A lot of future work has been planned: additional testing of the scalability of the model; integration of the model with pre-defined data bases and superficial learning; testing the model on the level of social schemas and interactions.

7. Dissemination

The results of the project has been disseminated at the annual conference of the Cognitive Science Society (Boston, 2011), the International Conference on Cognitive Modeling (Berlin, 2012) and the ReteCog Workshop on "Architectures of the Mind" (Madrid, 2011). We also presented intermediate results at the regular seminars of the Department of Cognitive Science and Psychology at New Bulgarian University.

References

- Gentner, D. (1983). Structure–mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155-170.
- Gentner, D. (1989). Mechanisms of analogical learning. In S. Vosniadou and A. Ortony, (Eds.), *Similarity and Analogical Reasoning*, 199-241. London: Cambridge University Press.
- Gentner, D., Loewenstein, J., Thompson, L., &Forbus, K. (2009). Reviving inert knowledge: Analogical abstraction supports relational retrieval of past events. *Cognitive Science*, 33, 1343-1382.
- Gick, M. L., &Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology* 12, 306—355
- Di Lollo, V. (2012). The feature-binding problem is an ill-posed problem. *Trends in Cognitive Sciences*, 16(6), 317–321.
- Doumas, L. A. A., Hummel, J. E., &Sandhofer, C. M. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review*, 115(1).
- Falkenhainer, B., Forbus, K.D., &Gentner, D. (1989). The Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence*, 41, 1-63
- French, R. M. (1999). Catastrophic Forgetting in Connectionist Networks. *Trends in Cognitive Sciences*, 3(4), 128-135.
- Holyoak K. &Thagard P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295-355.
- Hummel, J. &Holyoak, K. (1997). Distributed representation of structure: A theory of analogical access and mapping. *Psychological Review*, 104, 427-466.
- Kokinov, B. (1988). Associative memory-based reasoning: How to represent and retrieve cases. In T.O'Shea&V.Sgurev (Eds.), *Artificial intelligence III: Methodology, systems, applications* (pp. 51-58). Amsterdam: Elsevier Science Publ.
- Kokinov, B. (1994). The DUAL cognitive architecture: A hybrid multi-agent approach. In A. Cohn (Ed.) *European Conference on Artificial Intelligence*, (pp. 203-207). John Wiley & Sons.
- Kokinov, B., French, R. (2002). Computational Models of Analogy-Making. In: Lynn Nadel (ed.) *Encyclopedia of Cognitive Science*. Macmillan.
- Kokinov, B., Petrov, A. (2001). Integration of Memory and Reasoning in Analogy-Making: The AMBR Model. In: Gentner, D., Holyoak, K., Kokinov, B. (eds.) *The Analogical Mind: Perspectives from Cognitive Science*, Cambridge, MA: MIT Press.
- Kokinov, B., Vankov, I., Bliznashki, S. (2009). How Analogy Could Force Re-representation of the Target and Inhibition of the Alternative Interpretation. In: Kokinov, B., Holyoak, K., Gentner, D. (eds.). *New Frontiers in Analogy Research*. Sofia: NBU Press.
- Kuehne, S., Forbus, K., Gentner, D., & Quinn, B. (2000). SEQL - category learning as progressive abstraction using structure mapping. *Proceedings of the 22nd Annual Conference of the Cognitive Science Society*, (pp. 770–775). Mahwah, NJ: Lawrence Erlbaum.
- Loftus, E. F. & Palmer, J. C. (1974) Reconstruction of automobile destruction. *Journal of Verbal Learning and Verbal Behavior*, 13, 585-589.
- Petkov, G., Vankov, I., Kokinov, B. (2011). Deduction, Induction, and Analogy Modeled by the Same Mechanisms. *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*, Hillsdale, N.J.: Erlbaum.
- Petkov, G., Vankov, I., Kokinov, B. (2012) Dynamic Use of Multiple Analogies in the AMBR Model Causing Re-Representation of the Target. In N. Rußwinkel, U. Drewitz& H. van Rijn (eds.), *Proceedings of the 11th International Conference on Cognitive Modeling*, Berlin: Universitaetsverlag der TU Berlin.
- Snyder, M.; Cantor, N. (1979) "Testing hypotheses about other people: the use of historical knowledge", *Journal of Experimental Social Psychology* 15: 330–342.

- Von der Malsburg, C. (1986). Am I thinking assemblies? In G. Palm and A. Aertsen (Eds.), *Brain Theory*. Berlin: Springer.
- Wason, P. C. (1960) On the failure to eliminate hypotheses in a conceptual task. *The Quarterly Journal of Experimental Psychology*, 12, 129-140.
- Westen, D.; Blagov, P. S.; Harenski, K.; Kilts, C.; Hamann, S. (2006) Neural Bases of Motivated Reasoning: An fMRI Study of Emotional Constraints on Partisan Political Judgment in the 2004 U.S. Presidential Election. *Journal of Cognitive Neuroscience*, 18 (11), 1947–1958

Appendix A

The experiment is designed to test the prediction of the AMBR model that a rejected analogical base may play a role in highlighting specific aspects of the target that will improve the mapping between the target and another, more appropriate base.

Design

We performed a one-factorial between-group experiment. The independent variable was the presence or absence of a bridge analogical base. The dependent variables were participants' judgments of how similar the stories and some of their aspects are on a 7-point scale.

Procedure

Each participant received a sheet of paper with three short stories printed on them. The instruction was to read carefully all three stories and to prepare for answering some questions about them. There were no time limits for reading. Everybody worked alone, with the presence of the experimenter in the room only.

The participants from the control group received the stories "Terrorist", "Tsunami", and "Emigrant" (in this order); whereas the participants from the experimental group received "Terrorist", "Kamikaze", and "Emigrant" (see more about the stories in the section Stimuli below).

After that, the participants from both groups received another sheet of paper with eight statements on each. The instruction was to evaluate on a 7-point scale how much they agree with each of the statements. The last four statements were equal for both groups and concern the similarity between the "Terrorist" and "Emigrant" stories, as well the similarity between some of their aspects. The first four statements differed for both groups and concerned the similarity between the "Terrorist" and "Tsunami" or "Terrorist" and "Kamikaze" stories, respectively. The object of analysis was the answers to these four, identical for both groups, questions.

Stimuli

The four stories "Terrorist", "Kamikaze", "Tsunami" and "Emigrant" consisted of 120-170 words each. The first three stories were described as journalistic coverage, the fourth one – as a letter to a friend. The "Terrorist" coverage was about a lonely man who had crashed with a car-bomb in a market in New Jersey. The "kamikaze" report was about the grandson of a kamikaze, hero from the war. The grandson has been just nominated as an ambassador of Japan to US. The story of the tsunami (a control story for the participants in the control group) was about a Japanese farmer who had lost his business because of a tsunami. The "Immigrant" story was a letter from the wife of the immigrant to her friend.

The questionnaire consisted of eight statements. The first four statements differed between the two groups. For the control group they served for evaluating the similarity between the "Terrorist" and "Tsunami" stories; for the experimental group – between the "Terrorist" and "Kamikaze" stories,

respectively. People had to evaluate how similar they feel the stories as a whole; and the actions of the main protagonists; the motives for their actions, and the nature of the persons as a whole.

The second group of four questions served for evaluating the similarity between the “Terrorist” and “Immigrant” stories according to the same criteria. These four questions were the same for both groups and were the object of analysis.

Participants

42 students from New Bulgarian University participated in the experiment for course credits. They were randomly assigned to both groups. Twenty four (24) of them – in the control group; the other 18 – in the experimental group.

Results

The mean rating of how similar the stories about the terrorist and the kamikaze was 2.25 ($s = 1.23$) for the control group, and 3.83 ($s = 1.79$) for the experimental group. The difference was statistically significant: $t(40) = -3.40, p < .01$.

The differences between the three aspects of the stories, (whether the actions of the characters are similar; whether the motives for the actions of the characters are similar; whether the characters are similar in their nature) were not significant: $t(40) = -1.18, p = .24$; $t(40) = -.80, p = .43$; $t(40) = -1.03, p = .31$, respectively.

Thus, the difference of the ratings for the overall similarity cannot be explained by a simple assimilation effect. Looking at each aspect of the stories separately, the participants in both groups do not differ in their ratings. However, it seems that people in both groups weight the different aspects of the stories differently in the context of the third story. In other words, people weight the various aspects of the mapped stories differently depending on the context. This means that they have different representations of the target situation as predicted by the simulation.